

Datasets in ATLAS

ATLAS software meeting
Database session

David Adams

BNL

May 28, 2002

Contents

- Purpose
- Definitions
- Dataset types
- Dataset interface
- Dataset implementation
- Transformations
- Fundamental xforms
- Processing distributed data
- Conclusions

Purpose

Identify a collection of event data

- For physics analysis or
- Calibration/alignment or ...

Appropriate granularity for bookkeeping

- Or “virtual data”

Composite view for distributing data

Definitions

Event ID

- Identifier which uniquely labels each triggered beam crossing

EDO (event data object)

- Unit of event data (i.e. data associated with a particular beam crossing)

Content ID

- Identifier which labels the content of an EDO
- In StoreGate this is the EDO transient type and a string key



David Adams

BROOKHAVEN
ATLAS NATIONAL LABORATORY

ATLAS datasets

ATLAS SW week – DB session

May 28, 2002

4

Definitions (cont)

Event view

- Input for processing or analysis of an event
 - Don't want to see all data
- Collection of EDO's associated with a common beam crossing
 - I.e. with the same event ID
- Each EDO in the event view is labeled with its content ID
- View is unambiguous: it contains no more than one EDO for any content ID



David Adams

BROOKHAVEN
NATIONAL LABORATORY

ATLAS datasets

ATLAS SW week – DB session

May 28, 2002

5

Definitions (cont)

Complete event view

- View which includes the parents of all EDO's in the view

Consistent event view


- Subset of a complete event view
- If an EDO in the view has an ancestor whose content ID is included in the view, then the associated EDO must be the ancestor
- We probably want to require that views are consistent

Definitions (cont)

Figures


- The following figures illustrate the terms defined above
- Each figure represents an event view
 - Ovals are EDO's
 - Connecting lines indicate parentage
 - Labels indicate content ID
 - Color indicates the job (sequence of algorithms) that produced the EDO

Definitions (cont)



Two complete event views with the same content.

Definitions (cont)



Two incomplete and consistent event views with the same content.



David Adams

BROOKHAVEN
ATLAS NATIONAL LABORATORY

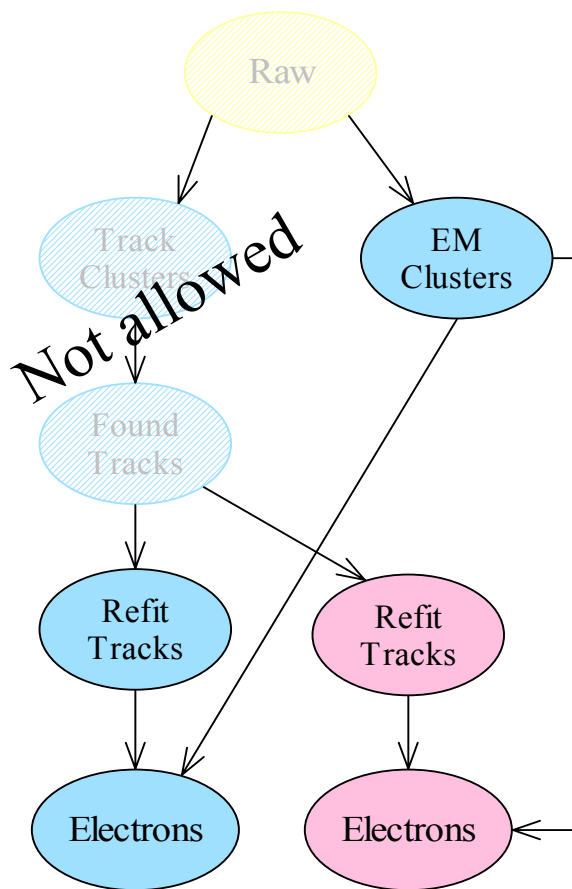
ATLAS datasets

ATLAS SW week – DB session

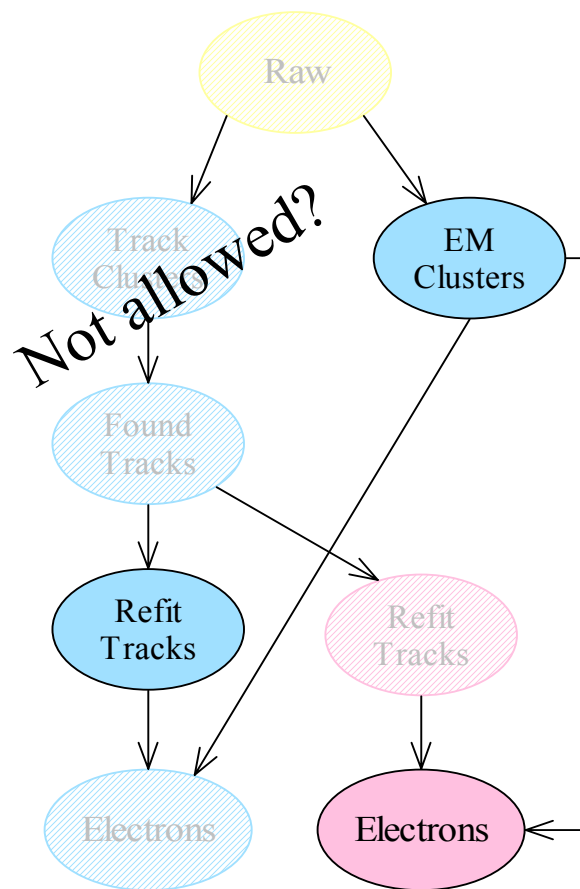
May 28, 2002

9

Definitions (cont)



Ambiguous event view.



Inconsistent event view.



David Adams

BROOKHAVEN
ATLAS NATIONAL LABORATORY

ATLAS datasets

ATLAS SW week – DB session

May 28, 2002

10

Dataset interface

Event range

- Collection of event ID's

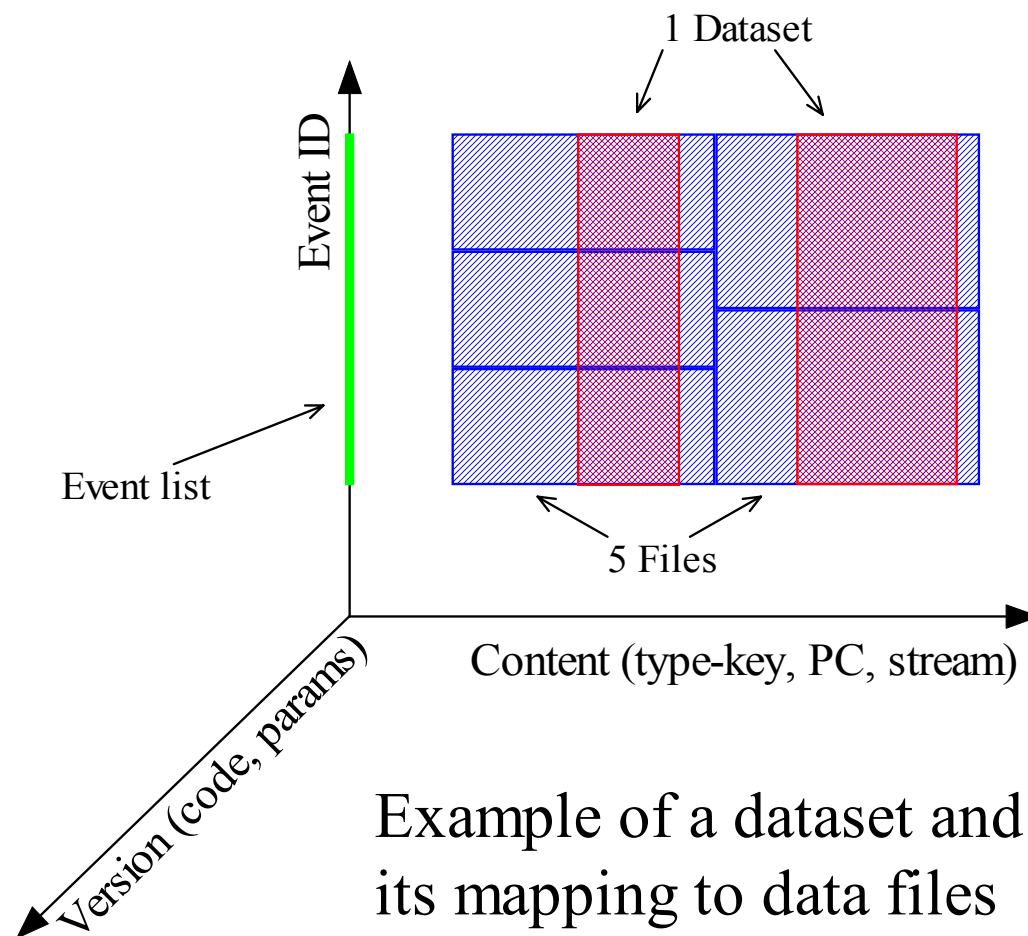
Content

- Collection of content ID's

Event data (event views)

- For each event ID-content ID pair:
 - A means to access the corresponding EDO or
 - A flag indicating the EDO is not included
- No other event data is included

Dataset interface (cont)



Dataset types

General comments

- The preceding interface may implemented in many ways
 - E.g. DB definition holding OID's for EDO's
- We would also like to have implementations that do not require DB access to locate data
 - User can process a small dataset on his or her laptop without network access
 - E.g. a file *is* a dataset
- Following describes some types of datasets

Dataset types (cont)

Physical dataset

- Knows a collection of files where its data may be found
- May know the files by logical or physical name
- Might restrict view
 - Which events to include
 - Which content is included

Dataset types (cont)

Composite dataset

- Holds references to other datasets and uses them to implement the interface
- Different types of compositions correspond to (some of) the fundamental xforms described later
 - Merge content
 - Merge events
 - Select content (content ID list)
 - Select events (event ID list)



David Adams

BROOKHAVEN
ATLAS NATIONAL LABORATORY

ATLAS datasets

ATLAS SW week – DB session

May 28, 2002

15

Dataset types (cont)

Virtual dataset

- Holds instructions for construction from unspecified input datasets and pointers to input dataset(s)
- In the language of “virtual data”, the instructions define a “transformation”
- Addition of the input datasets specifies a “derivation”
 - May also need to specify run-time environment

Dataset implementation

Datasets are used in many ways

- Inspection by humans
- I/O for processing in C++
 - And other languages
- Cataloging in DB's

Implementation

- Prefer something object oriented
- XML (?)

Transformations

Definition

- Datasets are produced by transformations which act on existing datasets (or other data)
- Datasets are the unit of data in a “virtual data” system

Fundamental xforms

- Any xform can be decomposed into these
- Intermediate states between these are appropriate for placing data (file, site, ...)

Fundamental xforms

0. Create (0-1)

- Create a dataset from a collection of EDO's
 - E.g. all raw data for a given time period

1. Add content (1-1)

- A sequence of algorithms is used to process each event individually
- New dataset includes the new EDO's
- Transformation is characterized by the expected input and output content

Fundamental xforms (cont)

2. Select content (1-1)

- Xform characterized by allowed content
- EDO's with content ID in this list are included in the output dataset

3. Merge content (N-1)

- Input datasets have the same event range
- Output dataset merges the contents of all input datasets
- If a content ID is shared by two input datasets, both must reference the same EDO



David Adams

BROOKHAVEN
ATLAS NATIONAL LABORATORY

ATLAS datasets

ATLAS SW week – DB session

May 28, 2002 20

Fundamental xforms (cont)


4. Event selection (1-1)

- Xform characterized by a test for accepting events (e.g. an algorithm)
- Output includes data for events that pass test


5. Event merge (N-1)

- Input datasets all have the same content and have non-overlapping event ranges
- Output has the same content and the union of the input event ranges


Fundamental xforms (cont)




Create content




Add content




Select content



Merge content



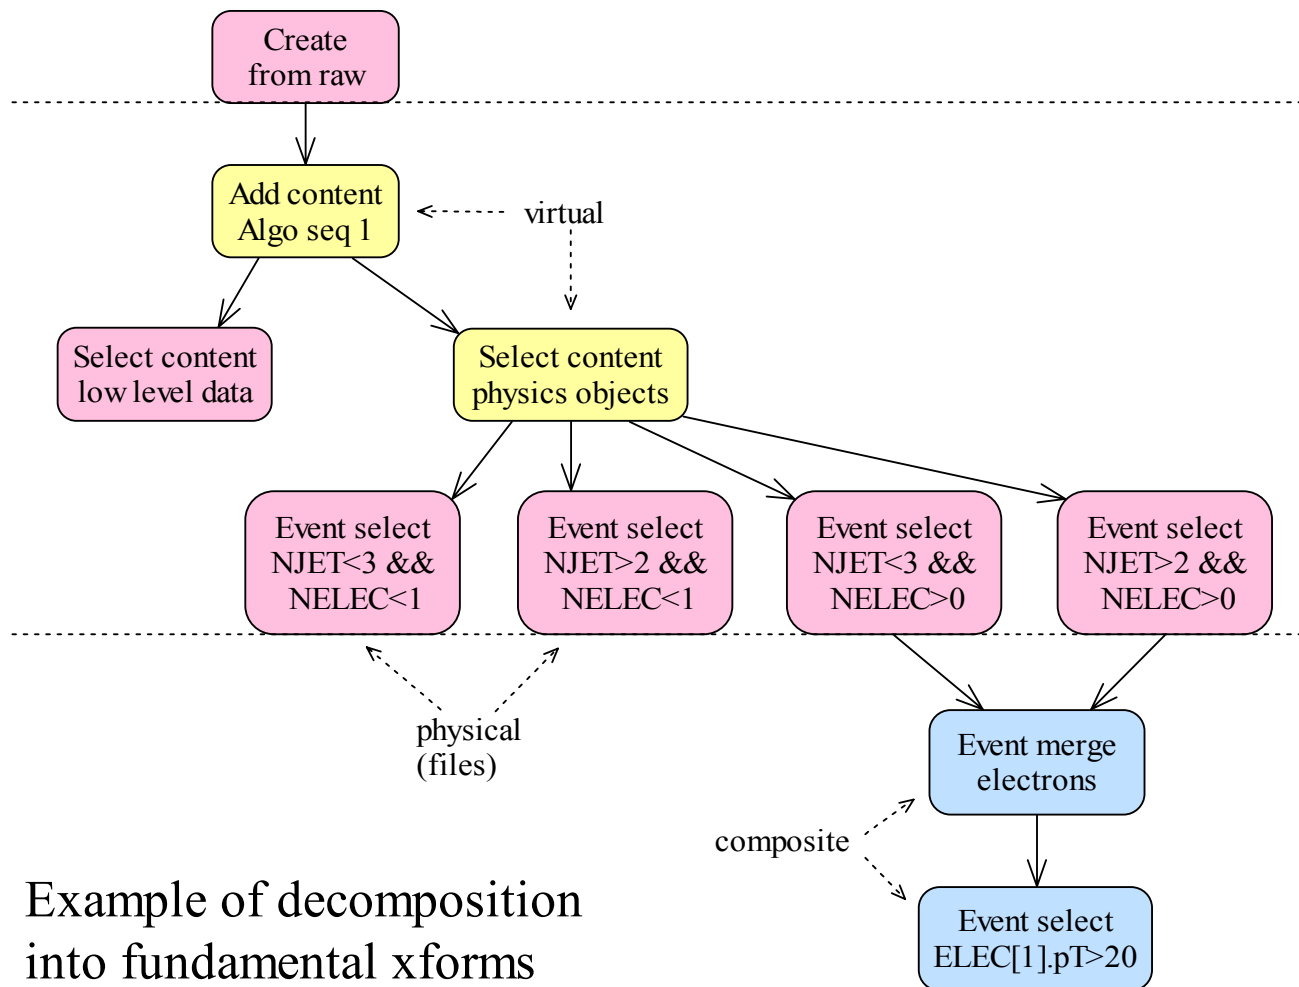
Event selection



Event merge

Allowed dataset transformations. In each case the blue input dataset or datasets are transformed into the red output dataset.

Fundamental xforms (cont)



Example of decomposition
into fundamental xforms



David Adams

BROOKHAVEN
ATLAS NATIONAL LABORATORY

ATLAS datasets

ATLAS SW week – DB session

May 28, 2002

23

Fundamental xforms (cont)

6. Representation (1-1)

- It is possible to have a two datasets which reference the same collection of EDO's through different collections of files
 - (Because of object replication)
 - Such datasets are equivalent
- Call this a representation xform
- Useful example is to copy only the selected events from the files in the input dataset

Processing distributed data

Decomposition

- Composite dataset can be decomposed into other composite or physical datasets
- Grid locates site(s) with data for each piece
- Scheduler selects an appropriate site
- Job(s) at that site produce new datasets

Composition

- Datasets (and other results) are gathered and assembled to create the final (likely composite) dataset



David Adams

BROOKHAVEN
ATLAS NATIONAL LABORATORY

ATLAS datasets

ATLAS SW week – DB session

May 28, 2002

25

Conclusions

Granularity

- Dataset is the appropriate granularity for processing, analyzing and sharing event data

Composition

- We can optimize disk access with data placement based on content and event selection
- Composite datasets can record this placement and provide transparent access to the data
 - With the help of GRID services
 - Assuming transformations decomposed similarly



David Adams

BROOKHAVEN
ATLAS NATIONAL LABORATORY

ATLAS datasets

ATLAS SW week – DB session

May 28, 2002

26

Conclusions (cont)

Is dataset the same as ADB event collection?

- Good question
- Similar properties
- ADB does not envision the types described here
 - Reference to files
 - Composite nature

Virtual data system

- Datasets are the (an) appropriate granularity for the unit of data in a VDS